California Enterprise Architecture Program
Office of the State CIO

# California

## *Service-Oriented Architecture (SOA) And Federated Identity Management Vision*

November 19, 2007

Draft

Lee Macklin, Acting Director
California Enterprise Architecture Program
Lee.Macklin@ceap.ca.gov

# Foreword by the State CIO

My Friends,

For three years now, we have been steadily working to establish a framework for enterprise-wide solutions and systems in California government. That framework includes major elements of our California State Information Technology Strategic Plan, the steps already taken to establish a statewide enterprise architecture, and the creation of IT governance bodies that encompass decision-making both from business and IT perspectives (i.e., the Enterprise Leadership Council, the Technology Services Board and the IT Council). The framework is largely in place (for more information on each of these topics, see www.cio.ca.gov).

Now, we must engage ourselves in the even more arduous task of actually building enterprise systems within the context of that framework. This will not be easy work. We have already teed up a broad portfolio of projects that will propel us towards our enterprise goals. For example, we have major infrastructure modernization efforts now underway at Department of Justice, Department of Motor Vehicles, Employment Development Department, Department of Corrections and Rehabilitation, Department of Health Care Services and Department of Technology Services. We are pursuing enterprise-wide business management systems through the Fi$Cal Project and in major ERP implementations underway at the State Controller's Office, Department of Corrections and Rehabilitation, Department of Transportation and the trial courts. We are building or rebuilding major case management systems all across government, including systems within Employment Development Department, Department of Motor Vehicles, Department of General Services, Department of Transportation, Department of Corrections and Rehabilitation, and the trial courts, as well as systems that deliver health and welfare benefits to millions of Californians.

But for these projects and systems to be truly "enterprise" in scope and operation -- as opposed to a series of siloed initiatives and efforts -- we must now commit ourselves to the cooperative development of enterprise standards that will meaningfully tie these systems together. The vehicle by which those standards will be developed is our enterprise architecture program. After prolonged consultation and deliberation, a consensus has developed that adopting a Service Oriented Architecture along with Federated Identity Management is the best course forward, the course most likely to lead to enterprise systems and capacity.

This document sets forth in the most basic terms possible our vision for that architecture and identity management approach. It also describes the nature of the issues that lie in the road before us and the type of decisions we will be making as we move down that road.

We seek your reactions, your support and your participation in this historic effort.

Clark Kelso
Chief Information Officer
State of California

# Executive Summary

It is time for an IT infrastructure in California that supports all the diverse lines of business from an *enterprise* perspective. It should be a standards-based environment that accommodates all stakeholders – different levels of government, as well as private industry partners. To build and maintain that environment over time, we will need to establish an inclusive governance model where all stakeholders can share their views and participate in creating standards.

There are two foundational components that make up this new IT infrastructure. First, an *Enterprise* Service-Oriented Architecture (SOA) infrastructure will host and manage new web services. Many of these business services will be implemented as shared web services. Second, an *Enterprise* Identity Management (IDM) system will manage all types of users in a consistent way. It will allow for various security policies to be applied as set by various security and privacy policy organizations.

The enterprise *environments* must be designed as *mission critical* since they will be hosting shared services. While some web services will likely be more "mission critical" than others, the enterprise SOA and Identity environments should be designed to accommodate very high availability and scalability. The consequences of critical services, such as identity provider services, not being available could be very significant as users would not be authenticated which would effectively cause many applications to appear as if they were not available. The unavailability of other services such as Verify SSN, Meds Eligibility, Vital Statistics, Verify Professional License, could also be disastrous.

The Roadmap section of this document lists the many steps that need to be accomplished to achieve success. It is anticipated that the roadmaps will be regularly updated resulting in many detailed projects. In order to accomplish the Roadmap, appropriate decision makers must provide strong policy on the project's sustainability, as well as identification of standards and processes and their enforcement. For example, to what degree will utilizing the Enterprise SOA and IDM infrastructures be mandatory? How will these Enterprise infrastructures be funded? How will they be sustained? Will standard language for SOA and IDM be enforced by the control agencies?

At some point, we may want to consider establishing an SOA Center of Excellence to manage and share policy, as well as serve as an exchange forum for all stakeholders. This approach is highly recommended by Gartner. In the meantime, several departments (DMV, EDD, DHCS, OSI, DSS, and DOJ) are initial members of the newly formed SOA Governance Group and they have agreed to work together to get the enterprise SOA and IDM infrastructures in place. Each of these departments has new projects underway which will utilize this shared, services-based infrastructure. Additionally, the Identity Management Workgroup was formed to lead the effort in defining federated identity management details.

As we move forward into the new world of SOA and federated identity management, we need to pay particular attention to the sustaining financial model. We want to encourage

stakeholders of all types to use this new enterprise environment. So, it is important that the financial model not be an impediment for providers and consumers of shared services.

## *California SOA and Identity Management Key Points*

Here are the key points of the California SOA and Identity Management environment:

1. Consistent with Federal standards and guidelines.
2. Designed to accommodate all stakeholders; it will support information sharing across government entities as well as public and private, subject to data sharing policies.
3. Users will be managed in a consistent way, and authentication will be at the level specified by the services they are accessing.
4. Extensive auditing capability of both user access details as well as the data they accessed.
5. Interoperability standards are defined; in most cases, communications will be via SOAP messaging, data will be formatted in XML, and services will have Web service interfaces. Note, SOAP was originally defined as Simple Object Access Protocol and later as Service Oriented Architecture Protocol. However, with SOAP version 1.2, both names were dropped and the official name is just SOAP.
6. This SOA environment is designed to handle online interactions via either voice or Web channels. Regardless of which input channel is used, the same services and interoperability processes will be invoked.
7. All users will be managed by a number of federated "identity providers" using standard identity protocols and interfaces. For example, "citizens" will be authenticated via a single identity service. All business service providers will "trust" this identity provider and not re-authenticate the user when accessing their services. State and local entities, as well as business partners could all be an identity provider for a certain class of users.
8. The preferred mechanism for formatting identity information is SAML (Security Access Markup Language) embedded in a SOAP message. Stakeholders may use either WS*, SAML, or CardSpace to share the identity information (subject to sharing policies).
9. Both "local" (departments) and "enterprise" environments are supported. That is, a department may choose to have its own SOA environment for use within the department. However, most shared services should be deployed to and managed by the enterprise environment running at DTS, Hawkins, and other major data centers.
10. Both the local and enterprise environment can be comprised of products from multiple vendors as long as they adhere to the interoperability requirements (SOAP, SAML, XML, Web service interfaces, etc). They will also need to meet appropriate scalability and availability requirements.
11. If all attributes for a given class of user are not located in one repository, then a *virtual directory service* will be used to provide a single or master view.
12. If a given user will have multiple accounts (and therefore, different and possibly conflicting information), then an *identity resolution* service will be used to determine that it is the same user across the different sources.
13. The degree of "opt-in" for "citizens" needs to be determined. That is, to what degree will a citizen have control over how and where their identity information is used?

14. The enterprise SOA environment primarily consists of: an enterprise service bus (ESB), a service registry, a module to govern web service policies, identity provider services, and operational policies and tools to manage the environment.
15. A service certification environment will need to be created as well as a governing group to manage shared services. They will determine who owns a service and the policies for modifying, extending, combining, or retiring a shared service. This will be very important as services are created and shared across state, local, and private entitles.

## *California SOA HIT Example*

Below is an example of how many healthcare stakeholders could utilize an enterprise SOA environment sponsored by the State. Since Health Information Technology (HIT) is a priority with this administration, the below drawing illustrates how new health-based services could be built to allow interoperability among the various industry providers resulting in new business services to constituents.

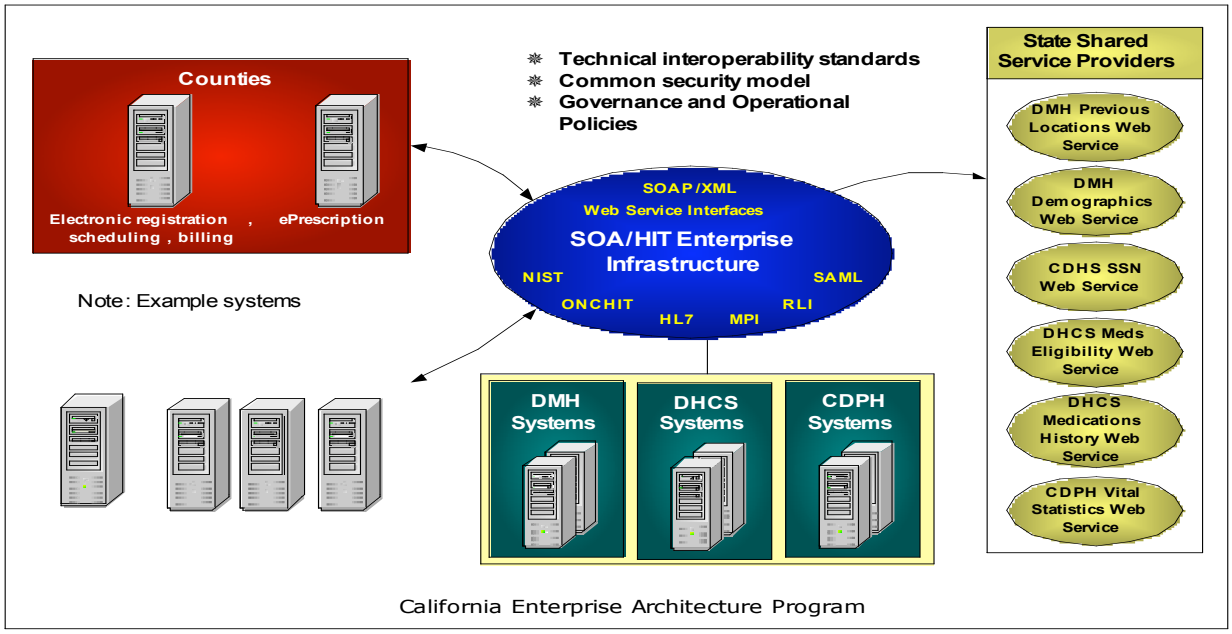

## Current Environment

In the past, most business applications were built in a "silo" fashion. That is, they were built for a specific purpose and they were "hard-wired" to other systems which made them difficult and expensive to change. Some of these systems have been modified to allow them to exchange information in more common formats (XML). However, for the most part the interfaces are still very platform and language dependent. There are very few standards set.

Additionally, there is no consistency of how user identities are managed. Policies vary even within the same department. This means a user must have multiple accounts where rules and policies may be applied differently. Thus it is difficult to manage changes, and worse, inconsistencies could result.

This problem will be magnified as more users move to conducting business online, particularly via the web but also via automated call centers. The State Portal has been completely redesigned and now provides a much better user experience. However, the infrastructure behind the portal also needs a major upgrade to an SOA environment.

## A Better Environment

As we migrate from a silo environment to a business services-based environment, we need a new infrastructure. The new business services will be implemented as web services – many will be shared services. For example, why not provide a single social security verification service and a single address verification service that all line of business applications can use instead of each building their own?

The new SOA environment provides opportunities to offer new business services that don't exist today. For example, the Department of Health Care Services has put into production a new set of shared web services that verify SSN and retrieve social security information, provide MEDS eligibility information, and provides a single view of medications history – integrating information from both government and private provider sources. All information is available real time, and the users access the information via standard web service interfaces (SOAP messages) and data is in XML format.

This does present some new challenges. First, we must have governance around the initial funding and project mechanisms for the SOA enterprise and IDM enterprise infrastructures, as well as a model for sustainability. Second, there must be governance around the management of the services. What is the process for deploying and certifying a service? What is the process for modifying a service? How will access to services be controlled? How do we deal with different security policies for different services? How do we authenticate different classes of users who have different authentication policies in a consistent way?

It is obvious that we must find ways to simplify and standardize user access to web services. Due to the complexity of California government, a federated approach will most likely be required to satisfy all the stakeholders.

Additionally, we need to put a stake in the ground regarding interoperability standards. Which interface and protocol standards will we choose? Which standards will be set for how data is exchanged?

This is an even bigger challenge when one considers that this impacts more than just state government. To achieve the biggest bang for the buck, different levels of government as well as partners in private industry must be included. So, a flexible standards model must be chosen.

The bottom line is we need to shift our thinking to an "Enterprise mentality" and leverage IT assets through reusable web services across lines of businesses and security domains.

This does not imply that we need to "pull the plug on the mainframe", or move from COBOL to .NET (for example). It does require that we agree on standard interfaces and standard policies regardless of how the services are actually implemented.

# Enterprise SOA Infrastructure

Will there be one SOA environment? No, there may be many local environments to handle department needs for services that are not shared outside the department. However, the enterprise SOA environments at major data centers, such as the Department of Technology Services and Department of Justice's Hawkins, should handle shared Web services.

So, what are the major components of an SOA environment? The primary components are a service bus, a service registry, a governance module to enforcing web services security, and the processes to manage the service environment. Other components are also important such as proxy servers, databases to store service and policy information, and application servers. We are referring to the components in the enterprise SOA environments as an *Enterprise* Service Bus and *Enterprise* Service Registry.

It is anticipated that at least half of the costs for the enterprise SOA infrastructure will be in consultant fees to identify, establish, and manage processes.

It is important that interoperability standards be set to ensure that the enterprise service bus and registry can communicate with other SOA environments that may reside at CalPERS, Franchise Tax Board, Counties, private industry partners, etc.

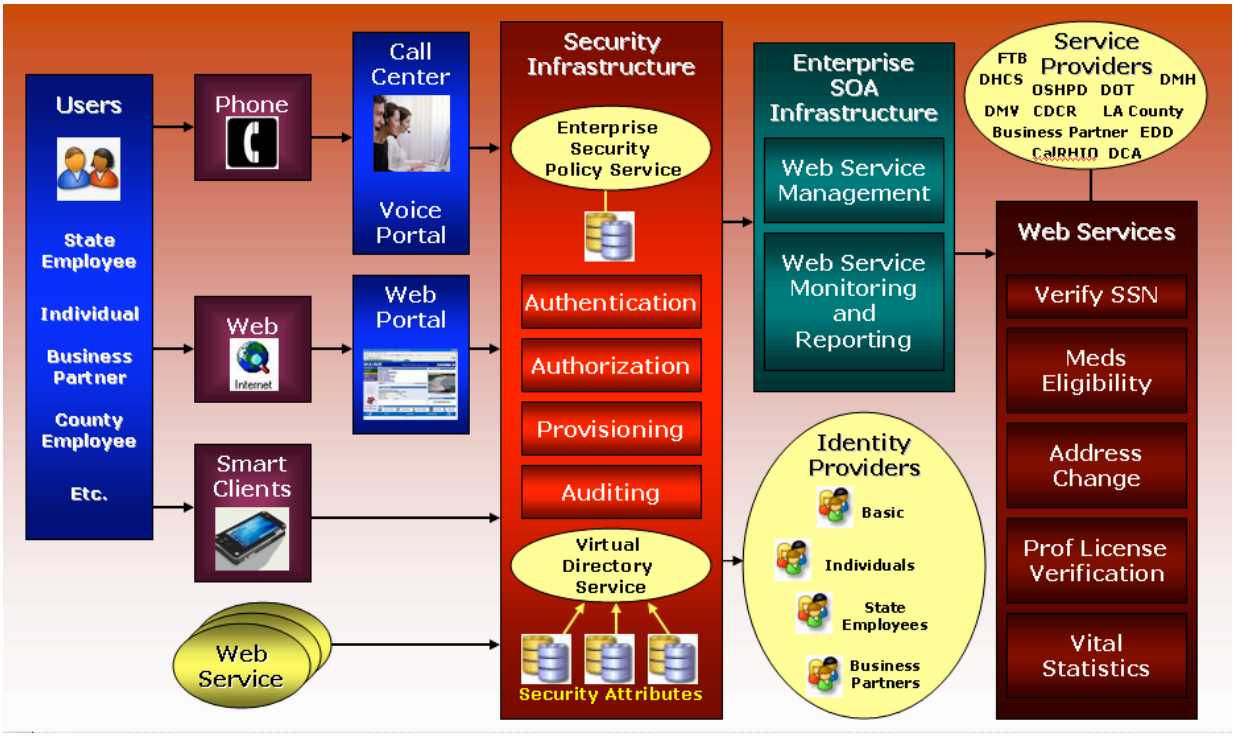


The above diagram illustrates how SOA and federated identity management work together. Regardless of how the user request was initiated (phone, web, smart client), the interaction

must first successfully pass through the security infrastructure before being routed to the appropriate web service.

Note, it is the combination of the Enterprise Service Bus and the Enterprise Service Registry that produces *transparency* of web services. That is, the users (consumers) of services should not need to know where the services are actually located. All consumer applications point to the ESB and it figures out (via consulting with the Enterprise Service Registry) where the services are actually located. This provides a lot of flexibility as the web services can be updated and moved without affecting the users.

Shared services come in two categories: *shared external* services are consumed outside of your department; *shared internal* services are consumed across projects or lines of business within your department or agency. A policy needs to be determined whether either or both of the shared service types will use the *enterprise* SOA service registry on a mandatory or optional basis.

In addition to the State portal, many departments will have line of business portals. They are free to choose the portal vendor and platform that the portal is built on (subject to State portal user interface, accessibility, and usability requirements as determined by the Office of eServices). However, if a portal is going to host shared external services, then a policy needs to be set as to whether or not these services will be registered with the *enterprise* service registry and access to these services managed by the *enterprise* identity management system.

Funding, spending authority, and project mechanisms must be determined for the major data centers to procure install, configure, and manage the initial Enterprise SOA infrastructure.

California SOA Reference Architecture

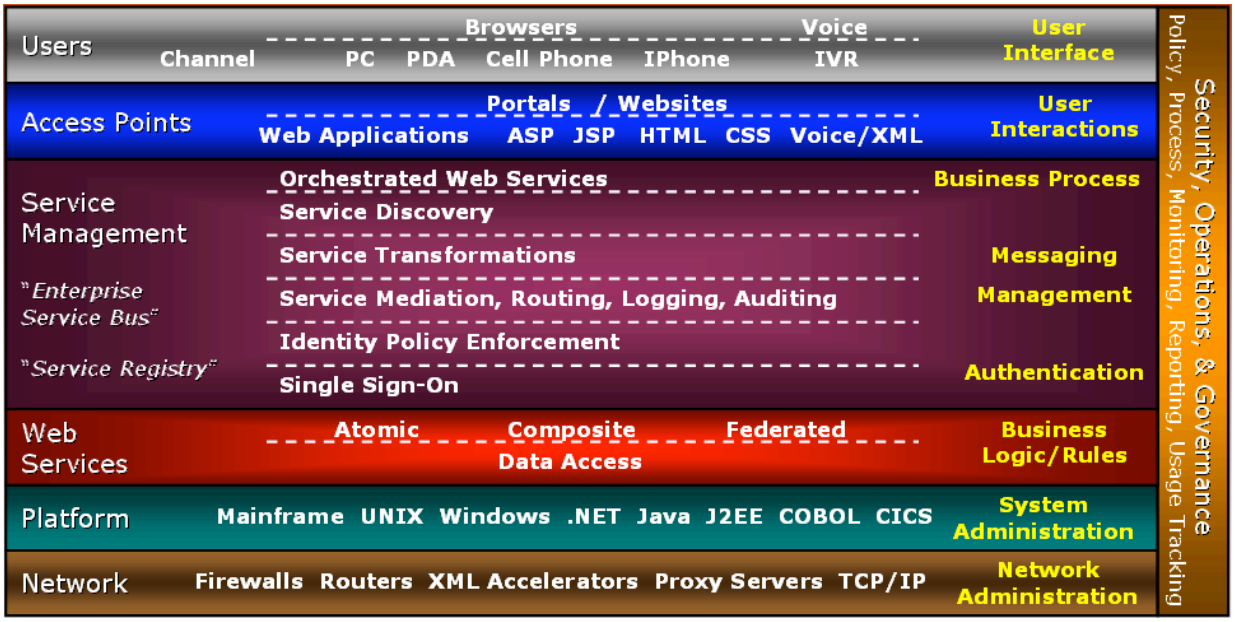

An SOA Reference Architecture is provided as a visual illustration to show how all the pieces fit into the enterprise SOA puzzle. It is important to note that the scope of this vision is limited to online services accessed via the web, voice through an IVR system or call center, or other smart clients. It does not address the broader physical security issues or paper processes.

A key part of the enterprise SOA infrastructure will be creating a certification environment separate from the production environment. New or modified web services can be tested to ensure that they play nicely before moving them to the production environment. Of course, the certification process must be defined and published so developers understand the environment and the process to get their new services certified.

## Enterprise Service Bus (ESB) and Service Registry

An ESB in conjunction with the service registry provides web services transparency. That is, all clients of a web service point to the ESB. Only the service registry knows the details (such as location and interface) of the web service and the ESB consults with the service registry to determine where to route the service request.

In addition to message routing, ESB's provide XML transformation, mediation, logging, and connectivity to all types of databases, external systems, and shared services.

Security policies should be established defining who can access as well as change information in the service registry. Products that are UDDI version 3 compliant provide this capability.

The ESB and service registry must have very high scalability and availability, as well as a bullet-proof recovery plan in case of a disaster.

## Web Service Providers

*Business services* are implemented as *web services*. They include the business logic, business rules, and data that make up the business functionality. They can be written in any language that supports web services (Java, .NET, etc.) but they must incorporate a web service interface. SOAP messaging is the standard mechanism for communicating with a web service. This hides the details of the language that the web service is written in. SOAP messaging is an industry standard that is language neutral, vendor neutral, and platform neutral.

Information about a web service (name, location, interface, and description) is formatted into a standard format called "WSDL" (Web Services Definition Language). The WSDL is registered with the service registry so the ESB can locate the web service.

Additionally, service providers may publish their security policy for accessing their service with the Security Policy Service. Identity providers may use the policy information to properly authenticate a user.

## Web Applications

Web *applications* are responsible for interactions with the user via a "session". They create and display the web pages that users see. They may "encode" security information into a link

that appears on a web page such as whether or not authentication is required. The web application is responsible for displaying the login page and formatting the user-provided information into a "SAML assertion". This is a special section in a SOAP message.
If Microsoft CardSpace is used, then the web application is responsible for displaying the list of available "cards", then passing the card to the identity provider.

## Medi-Cal Eligibility example

A business partner logs into their site and is presented with a link to check Medi-Cal eligibility. The identity service provided by the business partner queries the Security Policy Service to determine authentication requirements to access the DHCS Meds Eligibility Web Service. The identity provider performs the authentications and provides a SAML token with the appropriate information. The ESB consults the service registry to locate the DHCS Meds Eligibility Web Service and routes the SOAP message (with the embedded SAML token) to the policy enforcement point.

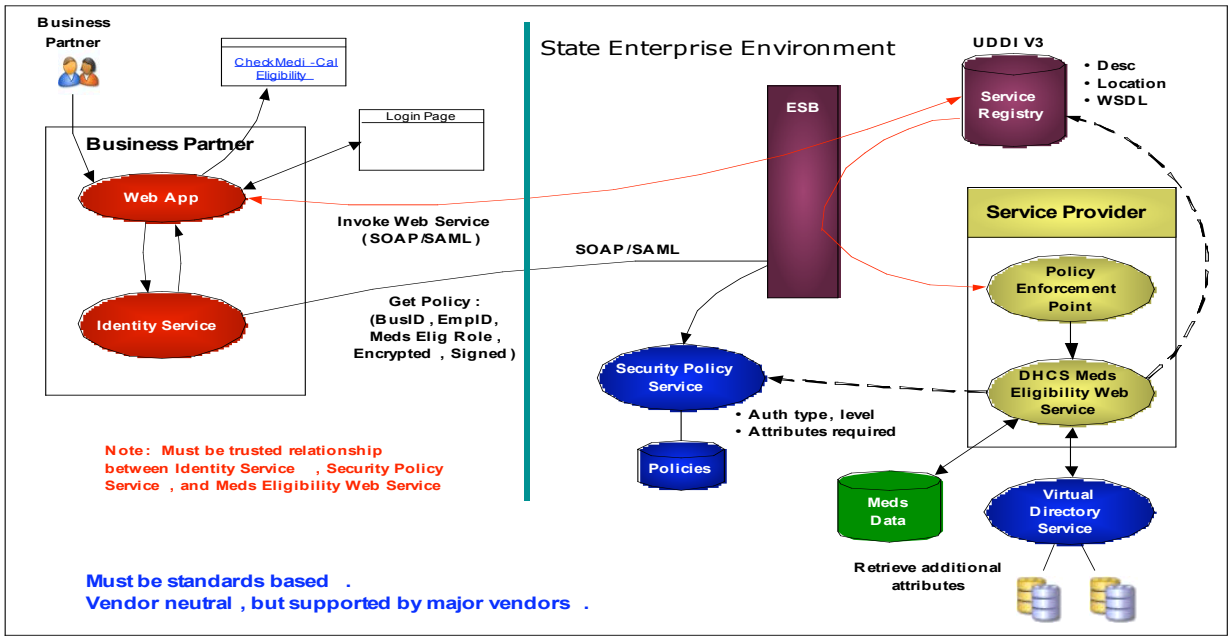

# Enterprise Identity Management

As we move to an online, shared services environment it is imperative that we implement a security system that ensures only authorized users get access to services that contain sensitive information. However, because security policies will vary widely across services, the identity management system must be flexible enough to determine which policy to apply, ensure that it is enforced, and meets auditing requirements. Today, we log into each system independently. This doesn't work in a shared Web services environment.

## *Federal Guide to Web Services Security*

In August 2007, the Federal government released the *Federal Guide to Web Services Security* (HIST 800-65). This is a detailed document that explains the many parts of federated identity management. The California SOA identity management model is compliant with this Federal guide. Included is this chart that shows the major standards grouped by function.

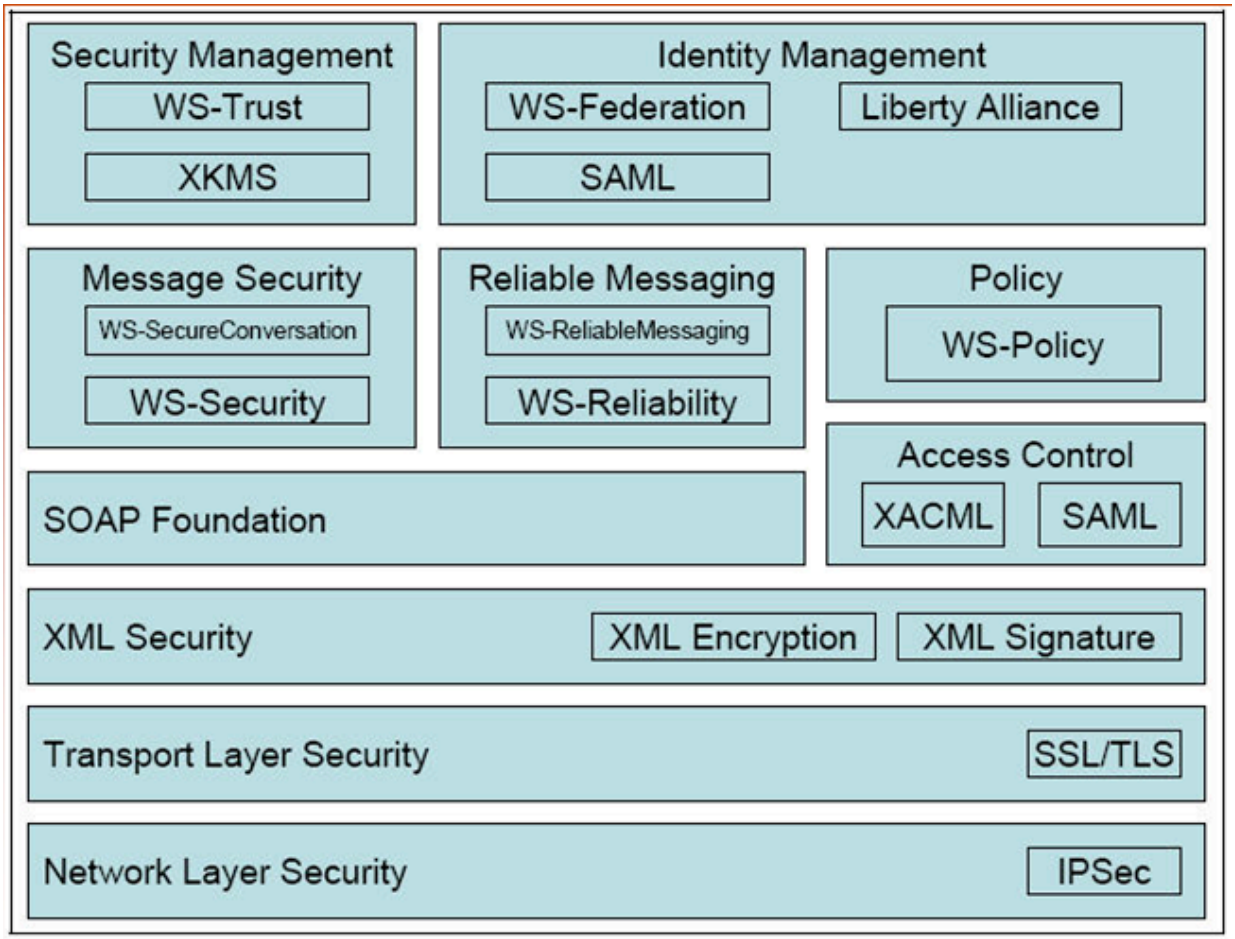


Federal Guide to Security Web Services (NIST 800-65 August 2007)

Key elements are:

1. XML Encryption (W3C standard) can be used to achieve Web service *message confidentiality*. This standard can be implemented using either SAML or X.509 certificates.

2. XML Signature (W3C standard) can be used to achieve Web service *message integrity*. This ensures that the message came from where you think it came from and was not hijacked (or altered) along the way. Again, either SAML or X.509 certificates can be used.

   *Note: The issue of whether or not digital signatures are valid in California government must be addressed. Specifically, we cannot accommodate Level 3 Authentication, as defined by Federal requirements (see below unless we can digitally sign a message). To clarify, we are referring to digitally signing a SOAP message, not signing a document (such as Microsoft Word, Adobe Acrobat, etc.).*

3. WS-Security, WS-Federation, SAML, WS-Trust, and WS-Policy (all OASIS standards) are the primary standards for managing authentication. Web Services Security (WS-Security) ensures *end-to-end* security instead of point-to-point (like SSL).

4. A service registry can be implemented using the Universal Description, Discovery, and Integration (OASIS) standard. This is where Web services are registered once deployed to the production environment. The enterprise service bus uses the service registry to locate the actual Web service and route a user's request to it.

5. All of the standards are implemented via XML sections (or "tags") inside a SOAP message.

## Authentication Levels

The Federal Electronic Authentication Guide (NIST 800-63) defines four levels of authentication as well as a risk assessment guide for mapping services to the appropriate authentication level:

1. Level One – Basic. User ID and Password via a simple password *challenge-response protocol* is the normal implementation at level one.

2. Level Two – Single Factor. Shared secrets (PIN, mother's maiden name, driver license number, etc.) are used to further verify a user's identity. Authentication is conducted by an *Identity Service Provider* via a secure authentication protocol, as defined in WS-Security (such as *SAML)*, and a security credential ("token") are required.

3. Level Three – Multi-factor. Authentication (via an Identity Service Provider) requires that the claimant prove through a secure authentication protocol that he or she controls the token (usually by a *digital signature*), and must first unlock the token with a password or biometric, or must also use a password in a secure authentication protocol to establish two factor authentication.

Tokens must be one of three types: "soft " cryptographic, "hard" cryptographic, or one-time password devices. A combination of SAML, XML Signature, XML Encryption, and X.509 certificates are often used to implement level three authentication.

4. Level Four – Only "hard" (physical hardware) cryptographic tokens are allowed. These are often "smart cards" that contain biometric information.

## *California Federated Identity Model*

In California, we have many different classes of users (citizens, non-citizens, state and county employees, business partners, etc.) who access government systems. Users need to be authenticated differently based on many different security policies which are set by the Web service owners. Therefore, a "one size fits all" identity management system is not practical.

Therefore, we intend to implement a federated identity management model. This means we will federate (or delegate) to a trusted party to handle the authentication based on a trust agreement. The goal is to establish multiple identity service providers, one for each class of user. An identity service provider will have the sole responsibility for authenticating its user class based on the security policy associated with the specific interaction (that is, the Web service they are trying to access) and the trust agreement among identity providers and Web service providers. The service providers will trust the identity providers and not re-authenticate the user. However, the service provider will inspect the credentials created by the identity provider for the user to ensure that this user has been properly identified and the credential actually came from the identity provider.

### Authentication Attributes

There are different classes of attributes that make up a user's identity.

1. Attributes that define "me" include name, address, date of birth, gender, fingerprint, birth certificate, etc.

2. Attributes that are *assigned* to me include user id, password, PIN, driver license number, social security number, employee number, account number, taxpayer id, Medi-Cal id, etc.

3. Identifiers assigned to my employer include employer id, federal employer identification number, etc.

4. Attributes may be combined into "master profiles" such as Individual Profile, State Employee profile, County employee profile, Incorporated Business profile, Professional Business profile, etc.

### Identity Service Providers

Identity providers perform authentication for a particular class of users (Individual, State Employees, etc.). If web service providers could agree on a single authentication policy for a given user class, then the identity provider would only need to authentication the user once. If

this is not the case, then the identity provider will need to first check the security policy and authenticate accordingly. If we only allow identity providers to access the security policy service, then keeping the number of identity providers low minimizes the risk of an inappropriate person accessing (and possibly altering) a security policy.

While there are several standards for authentication protocols and security tokens, SAML 2.0 is the preferred choice. The Secure Token Service (STS) will need to handle conversion if other types are accommodated.

The identity provider is responsible for creating the credential (security token). Upon successful authentication, it formats it as a SAML token and embeds it into a SOAP message.

A trust relationship must be established among identity providers and web service providers which is the definition of common understanding that they agree to.

### Security Policy Service

This is part of what is usually called SOA Governance. Some vendors combine (or bundle) security policy services with the service registry. Web service providers may establish their security policies and deploy them to the security policy service. They could include the authentication type and level as well as specific attributes that are required. Some products can also encrypt the attribute information providing further security.

One challenge is to restrict the number of users who have access to the security policy service. One way might be to only allow administrators in the service certification environment to change policy information. These administrators test new and modified web services before allowing them to be deployed into the production environment.

### Web Single Sign-On

This is the notion of signing in once and being able to perform multiple interactions without being re-authenticated. In practice this is difficult to achieve because different web services may have different security policies that require different levels of authentication. However, reducing the number of identity providers as well as common agreement among the web service providers and identity providers helps achieve this goal. There are specific protocols, such as Liberty Alliance and SAML 2.0 which are designed to handle web single sign-on.

Web service providers and identity providers generally form a “circle of trust”. The members agree to and implement a trust agreement.

## *Virtual Directory Service*

One item to be addressed is whether or not all attributes that identify a particular class of user can be stored in a single repository (LDAP, Active Directory, Tivoli, database, etc.). For example, for a State employee this can probably be achieved as the State Controllers Office and CalPERS have this information, so either could be the single identity provider for state employees. This means that all departments would federate to this single identity provider for authenticating any state employee regardless of the department they actually work for.

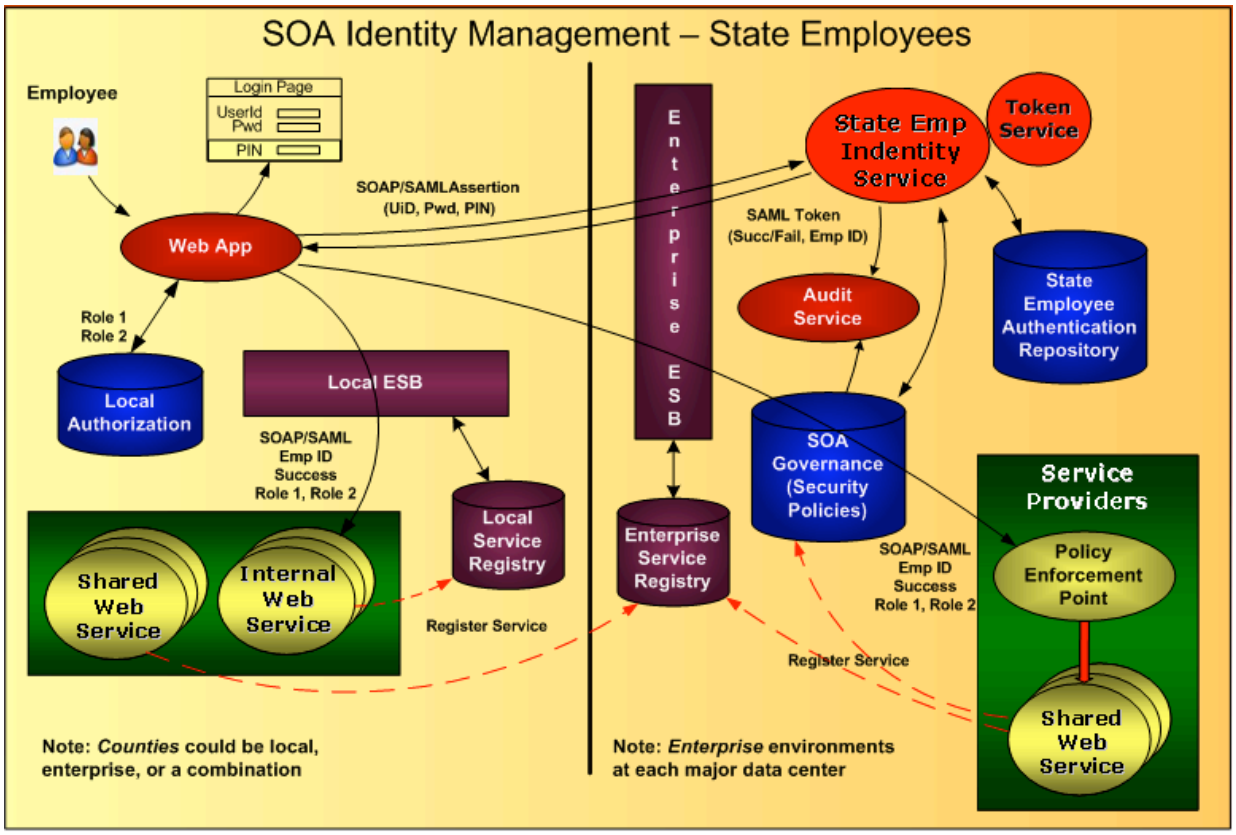


In the above example, a state employee logs into their department's web server application. It gathers the attribute information via a web form, formats the information in a "SAML assertion" (which is XML in a SOAP message), then forwards this to the identity provider service. If the majority of Web service providers can agree with the identity provider that a state employee will always be authenticated in the same way, then the identity provider simply verifies the information that the employee provided by comparing it to what is on file in the authentication repository. If successful, the identity provider creates a "SAML token" (a security credential), places it in XML format in a SOAP message and returns it to the Web application.

If we cannot agree on a single authentication process for state employees, then the identity provider must "look up" the policy in the SOA governance repository and authenticate according to this specific policy. On one hand, this is more flexible as it allows users to be authenticated differently. On the other hand, the user may need to be re-authenticated during the next interaction because that Web service has a different security policy. However, in either case all the attributes are stored in a single location. It is simply a choice of which ones to authenticate against.

For other classes of users, it is likely that user attributes may be managed and stored in different repositories. In this case, a *virtual directory service* should be used to obtain a "single view" of the user.

This might be the case with Individuals (citizens, non-citizens, visitors, etc.). For example Meds ID might be maintained by the Department of Health Care Services, Drivers License number by the Department of Motor Vehicles, and Adjusted Gross Income by the Franchise Tax Board. Different Web service security policies might specify that these user attributes be verified as part of the user authentication process. A virtual directory service makes this appear as if all the attributes were stored in one place.

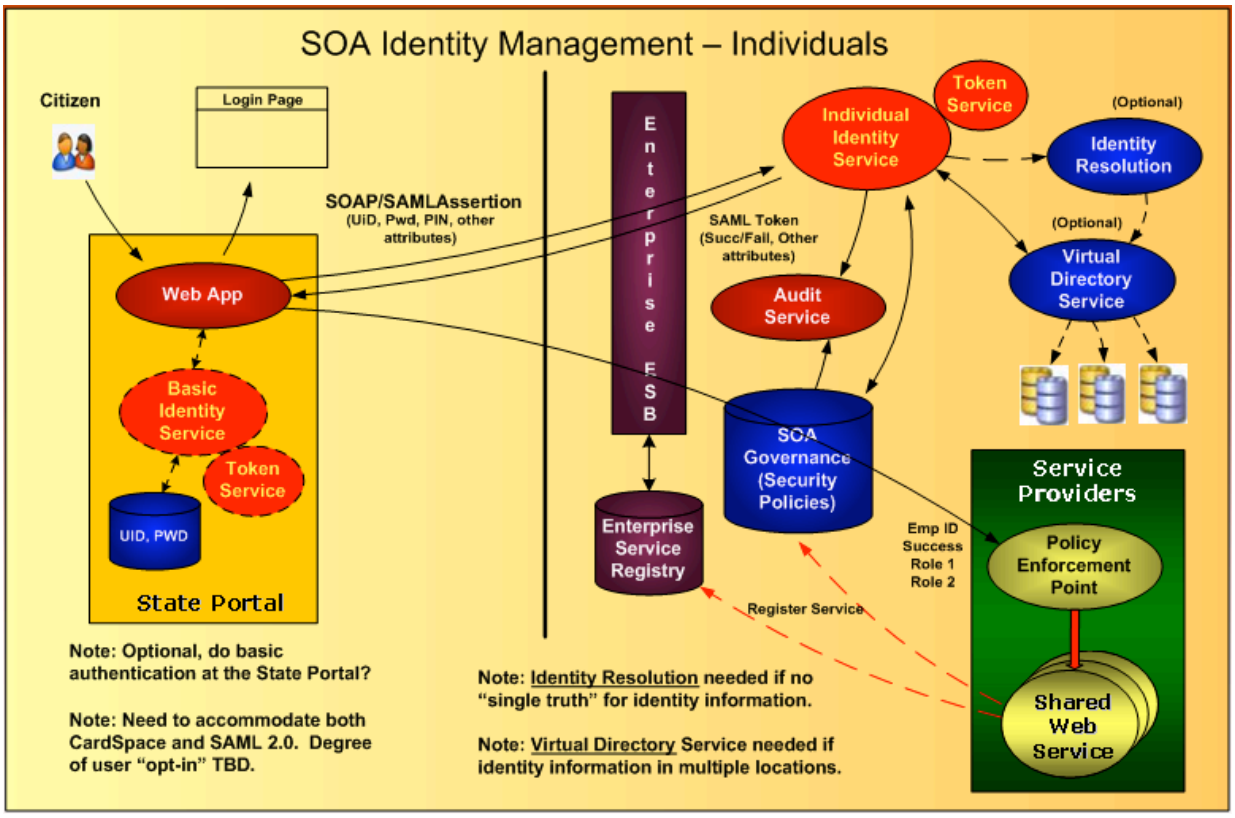


With regard to citizens, we may choose to provide "basic" (User ID, Pwd, PIN) authentication at the State Portal. This means that an identity provider beyond the portal would only need to be invoked if a higher level of authentication was required.

## *Identity Resolution Service*

If a user's information is defined in multiple locations – AND the information might be defined differently, then identity resolution software is a good solution for determining that it is the same user even though individual attributes might be different. This results because today a user must register with each system that they want to access. So, the name might be a little different (might use middle initial or nickname, or maiden name in some cases), the street address might not be spelled exactly the same (or abbreviated), and some attributes values might be null or missing. Here is a good example.

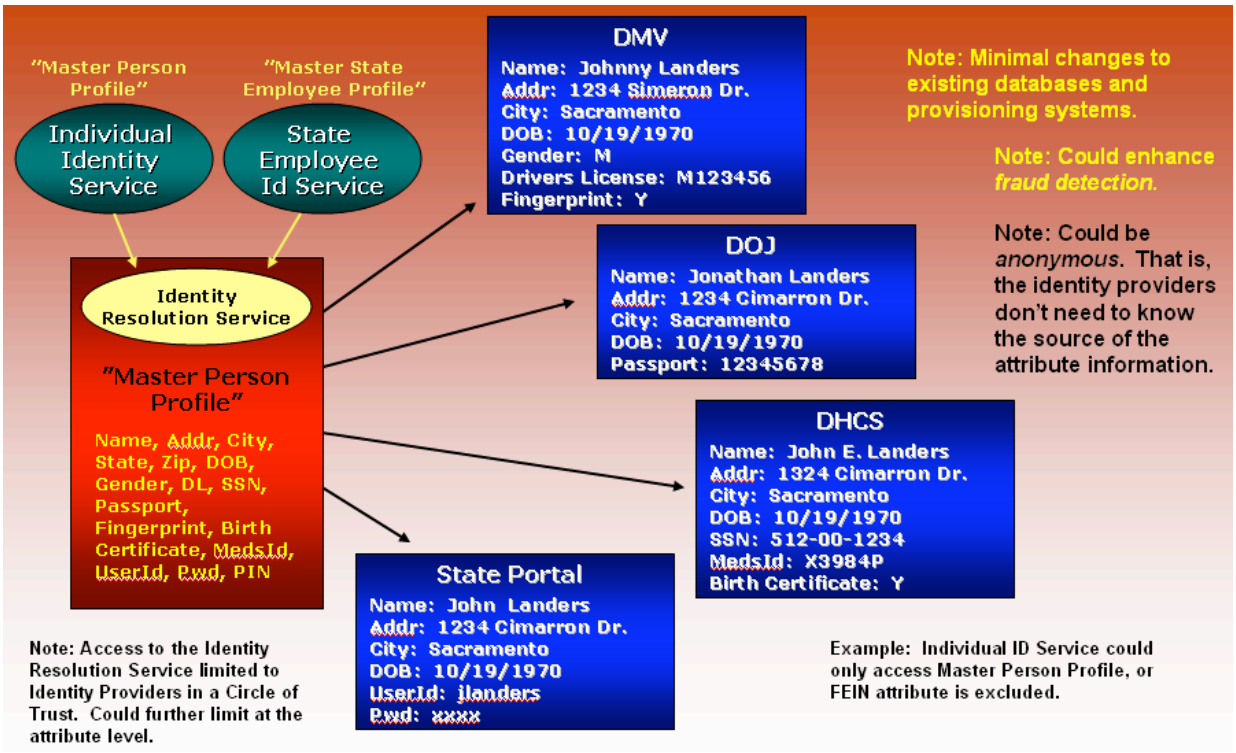


Identity Resolution

Is this the same John Landers? Probably, same city and date of birth, very similar name and addresses. A master profile is built that consist of "pointers" to the actual records. So, data is usually not replicated which provides flexibility in situations where the requestor cannot know the source of the data due to purpose of use or other restrictions.

Additionally, this is a great way to detect fraud since information can be compared across data sources. Data entry errors can also be caught as well as duplicates.

## *Change Address Example*

The following diagram illustrates a citizen updating his address. This could be a link on one of the state portal web pages.

1. The web application sees that authentication is required, displays the login page, extracts the userid, pwd, and pin that the user typed in, formats them into a SAML Assertion, creates the SOAP message, then sends it to the ESB.
2. The ESB looks up the Individual Identity Service in the Service Registry, then routes the SOAP message (with the SAML information embedded) to the identity service.
3. The Individual Identity Service looks up the security policy for the Address Web Service, sees all the attribute values are already in the SAML assertion, validates the information against the data on file (in the authentication repository).
4. If it matches, then the Individual Identity Service creates a SAML token ("credential") containing "success", wraps it in a SOAP message and sends it back to the ESB.

5. At this point, the ESB could inspect the message and see that the user has been properly authenticated, lookup the Address Web Service in the Service Registry, then route the SOAP message to the Policy Enforcement point that protects the Address Web service.
6. The Policy Enforcement point (proxy server, XML gateway, etc.) will do a final check of the credential before passing the message to the Address Web service.
7. If we wanted to notify other places that this user has changed his address, then we could call the Address Changed Notification web service which would send a message to the other data owners so they can update the address. This could be workflow enabled using a BPEL process.

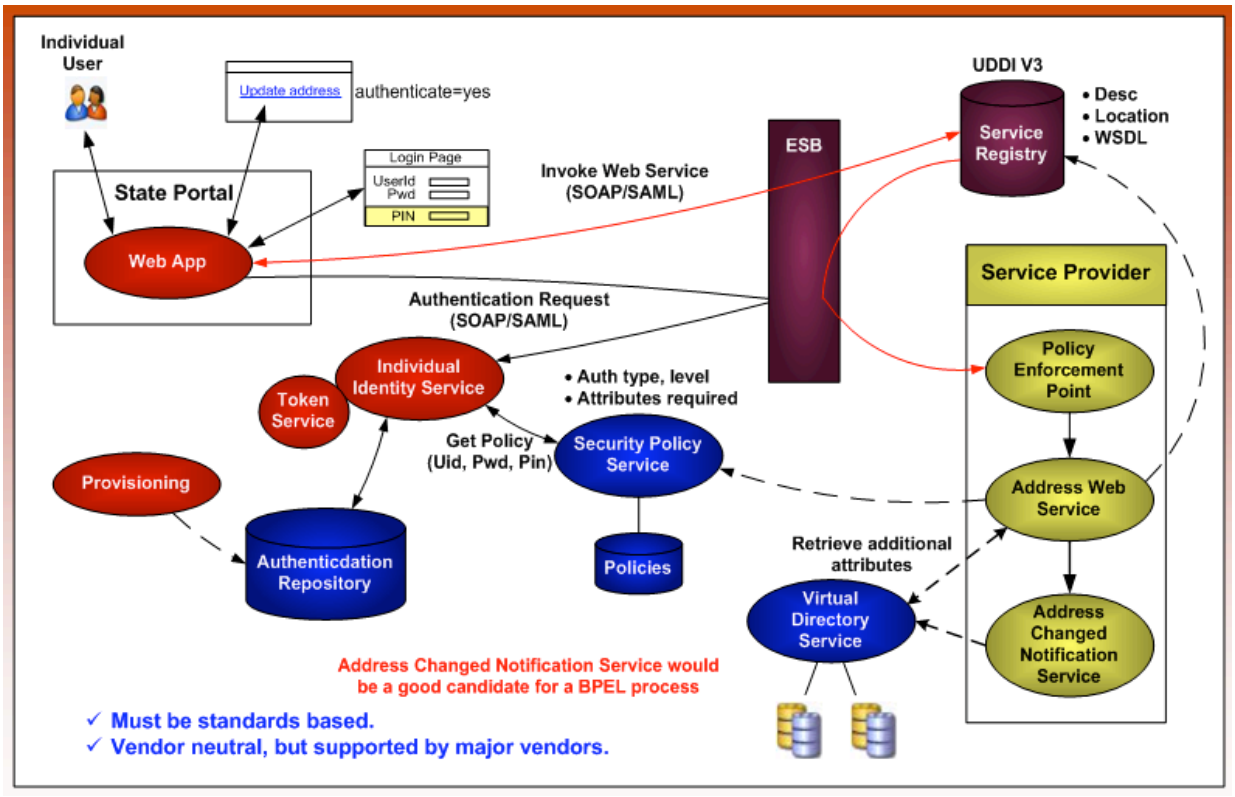


One of the big decisions to be made is how to define and where to store the security *policies*. There are several possibilities, but the current thinking is to use an SOA Governance module which serves as a Security Policy Service. Often, this module is combined with the Service Registry depending on vendor implementations.

This same challenge exists regarding how to best architect where the security *attributes* are stored and managed. These would typically be the roles that a user is a member of, but attributes are not limited to just role-based information. It is anticipated that these user attributes will be maintained in local repositories ("LDAPs"), but a "Virtual Directory Service" could "front-end" the local ones making them appear as one.

## *The Security/Authentication Process*

Web services communicate via industry standard SOAP messages. SOAP messages are a type of XML and are supported by practically every vendor. However, there are multiple choices for handling the security part of the SOAP message. The two leading standards are SAML and WS-SX. While there are several variations of these standards and under certain conditions they work together, for the purpose of this vision document we need to support both. These standards are evolving and hopefully, they will eventually merge into one.

Additionally, X.509 certificates are widely used for message authentication, message integrity, and message encryption. SAML, in conjunction with XML Encryption and XML Digital Signatures can accomplish the same thing. It is anticipated that California's SOA and Federated Identity Management infrastructure will need to support both X.509 certificates and SAML. A related issue is whether or not the State should provide a managed PKI (Public Key Infrastructure) service. This would enable a "single chain of authority" from local-state-federal government. The counties have asked the State to provide this service. The federal government is also encouraging this approach and they provide a bridge where state's can get certified.

One more potential problem needs to be addressed. The policy for digital signatures currently resides with the Secretary of State office. There are restrictions on digitally signing documents. However, it is unclear on whether or not this applies to digitally signing SOAP messages.

## *Trust Model*

Another decision area will be how do we define and implement a trust model? A trust model states the conditions that consumers and providers of shared services agree upon. This includes restrictions based on user profile or role as well as broader sharing rules among providers. The model also states liabilities that the parties agree to. A decision will need to be made whether a single trust agreement can be drafted that covers all the different scenarios (perhaps via appendices) or will multiple trust agreements be needed (perhaps one for government-to-government and a separate one for government-to-private industry).

The bottom line is the enterprise identity management system must ultimately accommodate all stakeholders (levels of government, public/private). The governance for this very broad model will very likely need to be revisited as more classes of users are added to the infrastructure.

For a detailed description of federated identity standards see
http://cio.ca.gov/calT/pdf/IDM_Standards.pdf

# Governance Model

Since Enterprise SOA and Federated Identity Management is a highly shared environment, it is critical that a governance model be defined and implemented to manage an environment that is acceptable to all stakeholders. Standards must be set and policies determined for how services will be created, certified, modified, and retired as well as how security will be defined and enforced. Additionally, initial and sustaining funding models and the mechanisms to implement enterprise infrastructure must be agreed upon.

California is a very complex government primarily based on a peer relationship model. Therefore, the governance model consists of many groups each with a specific focus. Here is a brief description of some of the groups and how they relate:

The **Enterprise Leadership Council** (ELC) is a top-level council that was formed to address enterprise issues. This includes providing executive support for ERP consolidation, service oriented architecture (SOA), and federated identity management efforts underway. Details of the ELC can be viewed at http://cio.ca.gov/stateIT/governance/leadership.html.

**Enterprise Process Advisory Committee** (EPAC) was formed specifically to manage the consolidation of the state's financial, procurement, and HR systems. Recently, their charter was expanded to include statewide enterprise system projects. EPAC is a working committee under ELC. EPAC details can be viewed at http://cio.ca.gov/stateIT/ITcouncil/committees/epac.html.

**Information Technology Council** (ITC) was formed to develop a statewide IT strategic plan and adopt enterprise-wide IT standards and policies. Membership is primarily state CIO's. There are approximately 10 committees, 2 sub-committees, and 8 working groups under the ITC. For more information on the ITC, see http://cio.ca.gov/stateIT/ITcouncil/index.html.

One of the ITC committees is Enterprise Architecture. This group is focused on developing a statewide business reference model (BRM) and data reference model (DRM).

**The California State Information Security Office** is currently part of the Office of Technology Review, Oversight, and Security (OTROS) group within the Department of Finance. One of the responsibilities is to set statewide security policy, both physical and computer online. As a result of SB90, this office will be moved to the new Office of Information Security and Privacy Protection effective January 2008. For details on the current office, see http://www.infosecurity.ca.gov/.

**Office of Privacy Protection (OPP)** is primarily focused on the protection of consumer's privacy information. The office is currently part of the Department of Consumer Affairs. As a result of SB90, this office will be moved to the new Office of Information Security and Privacy Protection effective January 2008. For information on the current office, see http://www.privacy.ca.gov/.

**Office of Information Security and Privacy Protection**
SB90 (Aug 2007) created a new Office of Information Security and Privacy Protection in the State and Consumer Services Agency effective January 2008. The existing "S" portion of OTROS, that is the State Information Security Office at DOF, and the Office of Privacy Protection will be moved into the newly created agency-level office.

**CalOHI** is a department within the Health and Human Services Agency that implements HIPAA (Health Insurance Portability and Accountability Act) and represents California in the federally sponsored nationwide HISPC (Health Information Security Privacy Collaboration) project. They are included in this document because they set privacy and security policy within the health line of businesses. The SOA and identity management infrastructure must support and enforce these policies. For more details see
http://www.ohi.ca.gov/state/calohi/ohiHome.jsp

**Privacy and Security Advisory Board (PSAB)**. This new board reports to the California Health and Human Services Agency (CHHS) and was recently formed as part of HISPC Phase II. There are four committees under this board: Privacy, Security, Legal, and Education. PSAB will make policy recommendations to CHHS. The SOA Governance Group will need to understand the HIT policies recommended by PSAB to ensure they can be properly managed within the SOA and Identity Management infrastructure.

**The California Enterprise Architecture Program** is part of the State Office of the CIO. It is responsible for creating the vision and blueprint for enterprise SOA, enterprise federated identity management (IDM), the Service Reference Model (SRM), and Technical Reference Model (TRM). CEAP provides reference architectures, researches and makes recommendations on SOA policy, management, and interoperability among levels of government and public/private stakeholders. CEAP works with all lines of business and rolls up individual requirements and issues into a cohesive enterprise. CEAP coordinates with most of the organizations described in this document.

The **SOA Governance Group** was formed to determine how to initially get the enterprise SOA and federated identity management infrastructure in place. This group deals with funding and process issues, and it reports to the ELC. CEAP and the Identity Management Workgroup frame issues and make recommendations to the SOA Governance Group. They will also provide policies for shared service lifecycle management. Additionally, they will make decisions on how to best implement federated identity management.

For more information on California Enterprise Architectures efforts see
http://cio.ca.gov/stateIT/enterpriseArch.html.

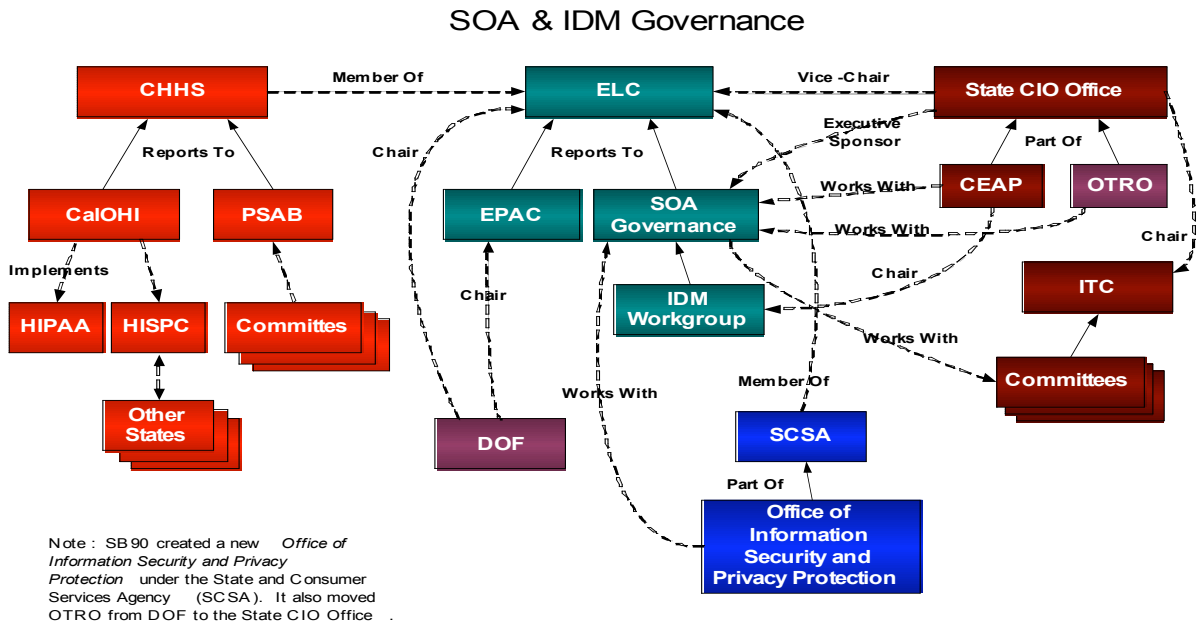


*Glossary of above Terms*:

| | |
|---|---|
| **CHHS**: | California Health and Human Services Agency |
| **CalOHI**: | California Office of HIPAA Implementation |
| **CEAP**: | California Enterprise Architecture Program |
| **DCA**: | Department of Consumer Affairs |
| **DOF**: | Department of Finance |
| **EPAC**: | Enterprise Process Advisory Committee |
| **ELC**: | Enterprise Leadership Council |
| **HIPAA**: | Health Information Portability and Accountability Act |
| **HISPC**: | Health Information Security and Privacy Collaboration (Federal) |
| **IDM Workgroup**: | Identity Management Workgroup |
| **ITC**: | Information Technology Council |
| **OTROS**: | Office of Technology Review, Oversight, and Security (until Jan 2008) |
| **OTRO:** | Office of Technology Review, Oversight (effective Jan 2008) |
| **PSAB:** | Privacy and Security Advisory Board |
| **SCSA**: | State and Consumer Services Agency |
| **SOA**: | Service Oriented Architecture |
| **State CIO**: | State Chief Information Officer |

# Enterprise SOA & IdM Roadmap

To implement the vision stated in this document, this roadmap lists the key activities and their approximate timeframes. It is anticipated that this roadmap will continually be updated as decisions are made and implementation progresses. More detailed timeframes will follow within each of the major areas.

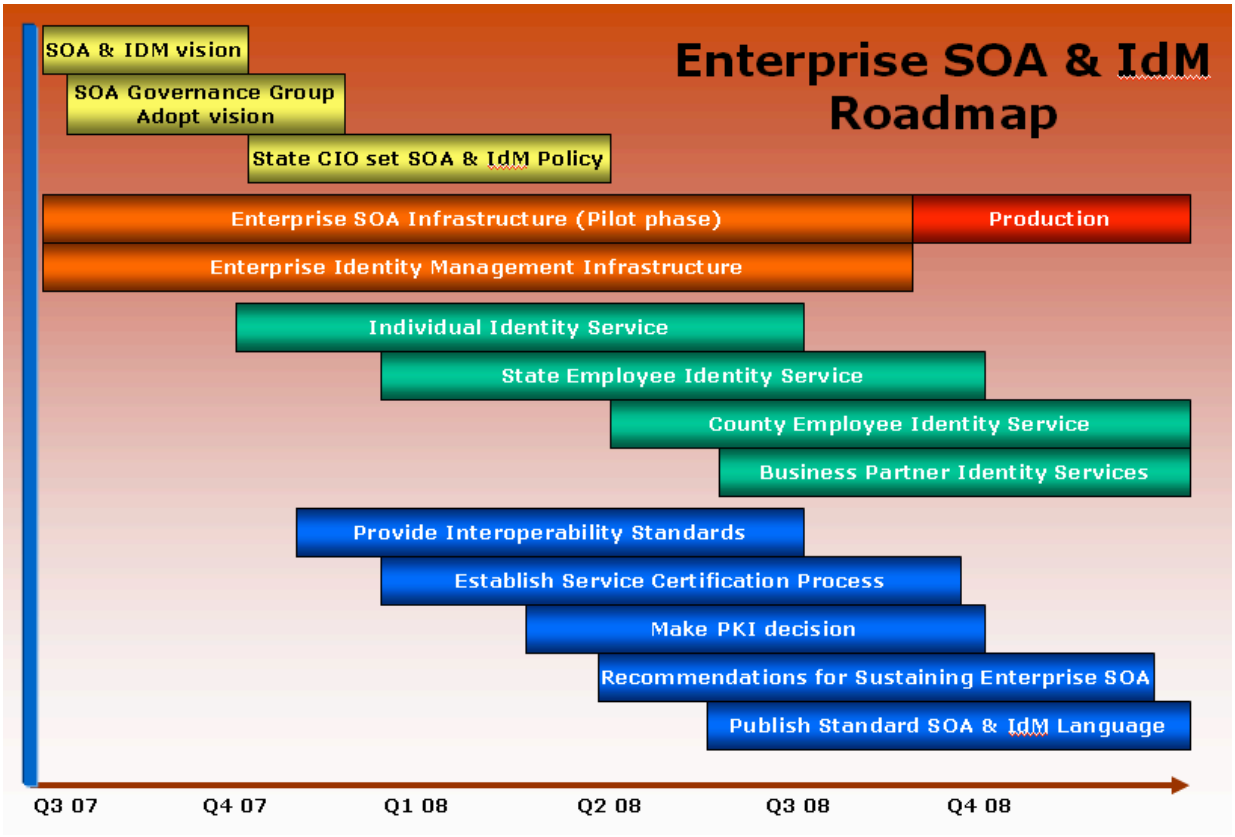


The general strategy is to first state the vision followed by several parallel phases. In the Initial Infrastructure phase, the SOA Governance Group will debate this vision followed by a recommendation that it be adopted as the general direction for moving to a new enterprise information technology infrastructure.

Next, we must sort out the details of how to fund and procure the enterprise SOA infrastructure. In parallel with the procurement effort, we must determine how to translate the EDD/DOL funding into a DTS procurement for an enterprise identity management (IDM) system.

Since we plan to migrate to a federated identity model, we must determine which user classes we plan to implement first. Based on initial discussions, it is recommended that we first define the Individual user and build the Individual Identity Provider Service. The Individual user could be a citizen, resident of California, legal alien, etc.

Assuming that X.509 certificates will be used in the federated identity process, as the Counties have requested a state-managed process, we need to make a decision on whether or not the state will provide a managed PKI solution. If the decision is to proceed, then a PKI program will have to be established. The State of Illinois has managed PKI in production and we could learn a lot and perhaps use their documents and processes as a starting point. We have had initial discussions and they have already provided some documents for our review. A managed PKI program includes establishing state-wide policies for all government stakeholders regarding how certificates will be structured, issued, and maintained as well as a certification policy and technical environment. Once we have sorted this out at the state level, we would then begin the process of becoming certified with the federal bridge. This would result in a single chain of authority from local to federal government meaning communication could occur at any level based on the trust policies and managed PKI certificate structure.

After we get the enterprise SOA & IdM infrastructures going, we need to focuses on creating documents that detail the standards and the processes for using the enterprise SOA and IDM infrastructure. These would be in the form of a series of recommendations provided by the SOA Governance Group to the Enterprise Leadership Council for adoption which would result in policy. This includes providing standard language that would be used in all funding request, project request, and procurement documents. Of course, the controlling agencies would need to also approve this language and agree to enforce it. This will provide vendors with a consist set of definitions so future services will be built in a standardized way.